Cisco Application Policy Infrastructure Controller OpenStack and Container Plugins Release 3.2(2), Release Notes

This document describes the features, caveats, and limitations for the Cisco Application Policy Infrastructure Controller (APIC) OpenStack and Container Plugins.

Cisco APIC OpenStack Plugins are used to deploy and operate OpenStack instances on a Cisco ACI fabric. It allows dynamic creation of networking constructs to be driven directly from OpenStack, while providing additional visibility and control from the Cisco APIC.

Cisco APIC CNI Plugin is used to deploy and operate Kubernetes clusters or OpenShift clusters on a Cisco ACI fabric. It allows dynamic creation of networking constructs to be driven directly from Kubernetes, while providing additional visibility and control from the Cisco APIC.

For the verified scalability limits (except the CLI limits), see the Verified Scalability Guide for this release. For the OpenStack, Kubernetes, OpenShift, and Coud Foundry Platform Scale Limits:

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of OpFlex hosts per leaf or vPC pair	40*	N/A
Number of endpoints per leaf or vPC pair	2000*	N/A

* same scalability values for OpenStack, Kubernetes, OpenShift, and Could Foundry Platform

For the CLI verified scalability limits, see the Cisco NX-OS Style Command-Line Interface Configuration Guide for this release.

Release notes are sometimes updated with new information about restrictions and caveats. See the following website for the most recent version of this document:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Table 1 shows the online change history for this document.

Table 1 Online History Change

Date	Description
July 10, 2018	Release 3.2(2) became available.
July 23, 2018	Added a new plugin for Juju Charms to include OpenStack Pike. For more information, see the New Software Features section.

Date	Description
July 30, 2018	Added a new plugin for Cloud Foundry and PCF for upstream compatibility. For more information, see the New Software Features section.
August 20, 2018	Cisco APIC OpenStack and Container Plugins, Release 3.2(2) is compatible with Cisco APIC, Release 3.2(3i) or later. For more information, see the Cisco APIC OpenStack, Container Plugins and Cisco APIC Compatibility Matrix and Changes In Behavior sections.
September 13, 2018	Added supported server platforms for the container solutions. For more information, see the Containers Guidelines section.
September 20, 2018	Added information about using external IPs for Load Balancing. For more information, see the Containers Guidelines section.
October 5, 2018	Added information about Cisco ACI OpenStack plugin works only with ceratin Neutron Red Hat package versions and the allowed address pair feature with the Cisco ACI plugin. For more information, see the OpenStack Guidelines section.

Contents

This document includes the following sections:

Cisco APIC OpenStack, Container Plugins and Cisco APIC Compatibility Matrix

Table 2 shows the Cisco APIC OpenStack, Container Plugins and Cisco APIC Compatibility Matrix

Cisco APIC OpenStack and Container Plugins	Cisco APIC	Supported	Guidelines and Restrictions
3.2(2)	3.2(2)	Yes	Cisco recommends that you upgrade the Cisco ACI fabric to 3.2(3).
3.2(2)	3.2(3)	Yes	Cisco recommends this set up. Note: When you upgrade, you must upgrade the Cisco APIC to 3.2(3) or later before upgrading the Cisco APIC OpenStack and container plugin to release 3.2(2). Otherwise, you may see instability on the attached top of rack (ToR) switch.

New and Changed Information

This section lists the new and changed features in this release and includes the following topics:

- New Software Features
- Changes In Behavior

New Software Features

The following are the new software features for this release:

Table 3 Software Features, Guidelines, and Restrictions

Feature	Description	Guidelines and Restrictions
Increase perf of opflex-proxy	Increase perf of opflex-proxy for larger deployments	None.
Cisco ACI supports OpenStack Pike with Juju Charms	Added a new plugin for Juju Charms to include OpenStack Pike. File name is: juju-charms-18.05-3.tar.gz For more information, see the Cisco Software Download website, click APIC OpenStack and Container Plugins and choose 3.2(2.20180710) release.	None.

Cisco ACI supports Cloud Foundry and Pivital Cloud Foundry	Added a new plugin for Cloud Foundry and Cloud Foundry and Pivital Cloud Foundry for upstream compatibility. File name is: dist-generics-3.2.2-cloudfoundry1.0-20180730.tar.gz For more information, see the Cisco Software Download website, click APIC OpenStack and Container Plugins and choose 3.2(2.20180710) release.	None.

Changes In Behavior

This section lists changes in behavior in this release.

- If you are going to upgrade, you must upgrade the Cisco ACI fabric first before upgrading the Cisco APIC OpenStack and container plugins. The only exception is for the Cisco ACI fabric releases that have been explicitly validated for this specific plugin version in the Cisco ACI Virtualization Compatibility Matrix.

  For more information, see the Cisco ACI Virtualization Compatibility Matrix at the following URL:

  https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/aci/virtualization/matrix/virtmatrix.html

- Starting in release 3.1(1) for OpenStack, the following changes were made to the unified plugin:
  - Adds support for the OpenStack Ocata release
  - Moves security group implementation from IPtables to OVS
  - Improves support for multiple OpenStack instances on the same CiscoAPIC cluster
- Security Groups for Opflex hosts are implemented natively in OVS, instead of using IPtables rules.

  If you are using an installer plugin distributed with this code, the appropriate configuration of Opflex hosts is automatically done. If you have your own installer, this change requires the following changes to the bridge configuration on all Opflex hosts:

  1. Create the br-fabric bridge, enter the following commands:

     ```
     # ovs-vsctl add-br br-fabric
     # ovs-vsctl set-fail-mode br-fabric secure
     ```

  2. Add a vxlan port to the br-fabric, enter the following command. The br-int_vxlan0 vxlan port on the br-int bridge is no longer needed and can be removed.

     ```
     # ovs-vsctl add-port br-fabric br-fab_vxlan0 -- set Interface br-fab_vxlan0 type=vxlan
     options:remote_ip=flow options:key=flow options:dst_port=8472
     ```

  3. Change the agent-ovs config file:

     ```
     "renderers": {

     "stitched-mode": {

     //"ovs-bridge-name": "br-int", <=== Remove this line.

     "int-bridge-name": "br-fabric", <=== Add this line.

     "access-bridge-name": "br-int", <=== Add this line.
     ```

```
"encap": {
          "vxlan" : {
              //"encap-iface": "br-int_vxlan0", <=== Change from br-int to br-fab.
              "encap-iface": "br-fab_vxlan0",
              "uplink-iface": "eth1.4093",
              "uplink-vlan": 4093,
              "remote-ip": "10.0.0.32",
              "remote-port": 8472
          }
      },
```

- Multiple OpenStack instances can share the same Cisco ACI fabric. Earlier versions of unified plugin would attach all OpenStack VMM domains to every OpenStack instance. This release allows cleaner separation by using this procedure:

  You must provision the VMM domains owned by each openstack instance using the new host-domain-mapping CLI command:

  ```
  # aimctl manager host-domain-mapping-v2-create [options] <host name> <domain name> <domain type>
  ```

  The host name can be a wildcard, which is indicated using an asterisk surrounded by double quotes ("*"). A wildcard means that the mapping should be used for all hosts. When more than one OpenStack instance shares the fabric, an entry must be created in this table for each VMM domain in use by that OpenStack instance. As an example, if one OpenStack instance is using VMM Domains "ostack1" and "ostack2", the following commands would be run on that OpenStack controller to put entries to this table:

  ```
  # aimctl manager host-domain-mapping-v2-create "*" ostack1 OpenStack
  # aimctl manager host-domain-mapping-v2-create "*" ostack2 OpenStack
  ```

  If the second OpenStack instance is using VMM Domain "ostack3", the following command would be run on that OpenStack controller to add an entry to its table:

  ```
  # aimctl manager host-domain-mapping-v2-create "*" ostack3 OpenStack
  ```

- Earlier versions only supported one logical uplink for hierarchical port binding or non-opflex VLAN network binding. In this release, you can have multiple links for those use-cases when using unified plugin.

  In order to use this feature, the AIM CLI has to be used to provide the mapping of physnets in OpenStack and an interface on a specific host. The following aimctl CLI command is used to configure this mapping:

  ```
  # aimctl manager host-link-network-label-create <host_name> <network_label> <interface_name>
  ```

  As an example, host h1.example.com is provisioned to map its eth1 interface to physnet1:

  ```
  # aimctl manager host-link-network-label-create h1.example.com physnet1 eth1
  ```

- Previously it was not possible for a single L3 Out to be shared across multiple OpenStack instances when using AIM, due to the fact that both OpenStack instances would attempt to use an External Network Endpoint Group of the same name. This release adds scoping of the Application Profile for the External Network Endpoint Group using the apic_system_id, which is configured in the `[DEFAULT]` section of the neutron configuration file.

- In earlier versions, the AIM plugin would take ownership of pre-existing L3 Outs when NAT was not being used, which led to scenarios where the AIM plugin would delete the pre-existing L3 Out in some corner cases. With this release, the AIM plugin will not take ownership of any pre-existing L3 Outs.
- Legacy plugin is not supported with the Ocata Plugins and will not be supported on future versions of OpenStack. The legacy plugin for Newton is supported. All customers are recommended to use unified mode for both Newton and Ocata.
- The OpFlex agent does not support client authentication. This means that the SSL certificate check must be disabled in Cisco APIC GUI.
    1. In the Cisco APIC GUI, on the menu bar, choose System > System Settings > Fabric Wide Setting.
    2. Ensure that the OpFlex Client Authentication check box is not checked.

Cisco ACI Virtualization Compatibility Matrix

For information about Cisco ACI, Kubernetes, OpenShift and OpenStack, see the *Cisco ACI Virtualization Compatibility Matrix* at the following URL:

https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/aci/virtualization/matrix/virtmatrix.html

Known Limitations

This section lists the known limitations.

- GBP and ML2 Unified Mode does not have feature parity with the ealier non-unified mode. In particular, it does not support the following features and for deployments that need some of these feautres, continue using the existing plugin configuration.
    - ESX hypervisor support
    - ASR1K edgeNAT support
    - GBP/NFP Service chaining
    - ML2 Network constraints
- Not all Unified mode features are support by the legacy plugin:
    - Support for OpenStack address scops
    - OpenStack address scopes are supported only in the Unified mode (where they are mapped to VRFs in the Unified model) and are not support in the earlier configurations.
    - Dual stack IPv6 deployment
- GBP and ML2 Unified Mode is a new mode of operation. So, while there can be a manual transition to this mode of usage, there is no automated upgrade from previous install to this mode.
- Dual-stack operation requires that all IPv4 and IPv6 subnets - both for internal and external networks - use the same VRF in Cisco ACI. The one exception to this is when separate external networks are used for IPv4 and IPv6 traffic. In that workflow, the IPv4 and IPv6 subnets used for internal networks plus the IPv6 subnets used

for external networks all belong to one VRF, while the subnets for the IPv4 external network belong to a different VRF. IPv4 NAT can then be used for external networking.

- Mirantis Fuel based plugins are not supported. For Ubuntu based installs, use the released Juju based installer.
- Cisco APIC OpenStack plugins do not support the Cisco ACI Multi-Site at this time.

Usage Guidelines

OpenStack Guidelines

- Currently the Cisco ACI OpenStack plugin works only with the Neutron Red Hat package versions up to and including 9.4.1-12.

  If you are using a version mentioned above the workaround is to set 'root_helper_daemon =' in the neutron.conf file and restart the neutron-opflex-agent. However, this is not recommended for production environments. Consult with Red Hat on how to make this change permanent in a OpenStack Platform Director (OSPD) environment.

- When using the allowed address pair feature with the Cisco ACI plugin, be aware of the following differences from upstream implementation:

  - The Cisco ACI plugin only supports the host (/32) specification for allowed_address_pair, not the CIDR/subnet specification.
  - As OpenStack allows the same allowed_address_pair to be configured on multiple interfaces for HA, the OpFlex agent requires that the specific VNIC that currently owns a specific allowed_address_pair to assert that address ownership using GRAT ARP.
  - When using the promiscuous mode, the vSwitch stops enforcing the port security check. To get reverse traffic for a different IP or MAC address, you still need to use the allowed-address-pair feature. If you are running tempest, you will see test_port_security_macspoofing_port fail in scenario testing, as that test does not use the allowed-address-pair feature.

- If you are using SLAAC, add a security group rule to allow ICMPv6 to the effected Neutron networks. For example, the following security group (ipv6-sg) allows the required traffic:

```
# openstack security group rule create --ethertype IPv6 --ingress --protocol 58 --src-ip ::/0
\ ipv6-sg
```

- Before performing an upgrade from 3.1(1) using OpenStack Director or attempting a Cisco APIC ID recovery procedure, all AIM processes on all controllers need to be shutdown. To shutdown all the AIM processes on all controllers, run the following command on the undercloud:

```
for IP in $(nova list | grep ACTIVE | sed 's/.*ctlplane=//' | sed 's/ |//') ; do

ssh -o UserKnownHostsFile=/dev/null -o StrictHostKeyChecking=no heat-admin@$IP \

"sudo systemctl stop aim-event-service-rpc; sudo systemctl stop aim-aid; sudo systemctl stop
aim-event-service-polling" ;

done
```

If upgrading, you do not need to explicitly restart the AIM processes as the upgrade will automatically restart them.

If attempting a Cisco APIC ID recovery, you must restart the AIM processes on all the controllers manually after ID Recovery is complete.

- Keystone configuration update

  When the OpenStack plugin is installed in the unified mode, the Cisco installer adds the required configuration for keystone integration with AIM. When not using unified mode, or when using your own installer, the configuration section must be provisioned manually:

```
[apic_aim_auth]

auth_plugin=v3password

auth_url=http://<IP Address of controller>:35357/v3

username=admin

password=<admin_password>

user_domain_name=default

project_domain_name=default

project_name=admin
```

- When using optimized DHCP, the DHCP lease times are set by the configuration variable apic_optimized_dhcp_lease_time under the `[ml2_apic_aim]` section.
  - This requires a restart of neutron-server to take effect
  - If this value is updated, existing instances will continue using the old lease time, provided their neutron port is not changed (e.g. rebooting the instance would trigger a port change, and cause it to get the updated lease time). New instances will however use the updated lease time.
- In upstream Neutron, the "advertise_mtu" option has been removed.

  Since the aim_mapping driver still uses this configuration, the original configuration which appeared in the default section should be moved to the aim_mapping section. For example:

```
[aim_mapping]

advertise_mtu = True
```

  It is set to True by default in the code (if not explicitly specified in the config file).
- GBP and ML2 Unified Mode allows coexistence of those OpenStack networking APIs on the same OpenStack and Cisco ACI instance, but they need to be running on different VRFs. This is a constraint that we may remove in future, but at this time this is the supported configuration.
- Unified mode has features not supported by the legacy plugin:
  - Support for Openstack address scopes and subnetpools

    OpenStack address scopes and subnetpools are supported only in the Unified mode (where they are mapped to VRFs in the unified model) and are not supported in the earlier configurations.

  - Dual stack IPv6 deployment

- If a default VRF is implicitly created for a tenant in ML2, it is not implicitly deleted until the tenant is deleted (even if it not being used anymore).
- Unified model impact of the transaction Model Updates in Newton.

  When GBP and ML2 co-exist, GBP implicitly created some neutron resources. In Newton, the neutron transaction model has been updated and has added a lot of checks. Some of those checks spuriously see this nested transaction usage as an error and log and raise an exception. The exception is handled correctly by GBP and there is no functional impact but unfortunately the neutron code also logs some exceptions in neutron log file – leading to the impression that the action had failed.

  While most such exceptions are logged at the DEBUG level, occasionally you might see some exceptions being logged at the ERROR level. If such an exception log is followed by a log message which indicates that the operation is being retried, the exception is being handled correctly. One such example is the following:

  Delete of policy-target on a policy-target-group associated to a network-service-policy could raise this exception:

```
2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource […] delete failed

2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource Traceback …:

2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource   File "/usr/lib/python2.7/site-packages/neutron/api/v2/resource.py", line 84, …

...

2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource     raise …

2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource ResourceClosedError: This transaction is closed
```

  Note: We are working with the upstream community for further support on Error level logs.

- When a Layer 2 policy is deleted in GBP, some implicit artifacts related to it may not be deleted (resulting in unused BDs/subnets on Cisco APIC). If you hit that situation, the workaround is to create a new empty Layer 2 policy in the same context and delete it.
- The ASR1K for edge-NAT does not provide support to create instances attached to external networks.
- If you use tempest to validate OpenStack, the following tests are expected to fail and can be ignored:

```
tempest.scenario.test_network_basic_ops.TestNetworkBasicOps.test_update_router_admin_state
```

- Neutron-server logs may show the following message when DEBUG level is enabled:

```
Timed out waiting for RPC response: Timeout while waiting on RPC response - topic: "<unknown>", RPC method: "<unknown>" info: "<unknown>"
```

  This message can be ignored.

- High Availability LBaaSv2 is not supported.
- OpenStack Newton is the last version to support non-unified plugin. OpenStack Ocata and future releases will only be supported with the unified plugin.

Containers Guidelines

- For OpenShift, the external IP used for the LoadBalancer service type is automatically chosen from the subnet pool specified in the ingressIPNetworkCIDR configuration in the /etc/origin/master/master-config.yaml file. This subnet should match the extern_dynamic property configured in the input file provided to acc_provision script. If a specific IP is desired from this subnet pool, it can be assigned to the "loadBalancerIP" property in the LoadBalancer service spec. For more details refer to OpenShift documentation here:

  https://docs.openshift.com/container-platform/3.9/admin_guide/tcp_ingress_external_ports.html#unique-external-ips-ingress-traffic-configure-cluster

  NOTE: The extern_static subnet configuration in the acc_provision's input is not used for OpenShift.

- You should be familiar with installing and using Kubernetes or OpenShift. The CNI plugin (and the corresponding deployment file) is provided to enable networking for an existing installer such as kubeadm or kargo. Cisco ACI does not provide the Kubernetes or Openshift installer.

- The released images for this version are available on dockerhub under user Noiro. A copy of those container images and the RPM/DEB packages for support tools (acc-provision and acikubectl) are also published on CCO.

- OpenShift has a tighter security model by default and many off the shelf Kubernetes applications such as guestbook may not run on OpenShift (if, for example, they run as root or open privileged ports like 80).

  Please refer to the following for details:

  https://blog.openshift.com/getting-any-docker-image-running-in-your-own-openshift-cluster/

- When running OpenShift, 'oc new-app' tries to reach Github. If you are running behind a proxy, due to OpenShift's issues with handling the proxy environment variable when it is set on the compute node, this connection may fail. This is an Openshift issue and not related to the networking provided by Cisco ACI.

- The supported server platforms for the container solutions are:

  — Kubernetes on Ubuntu 16.04

  — OpenShift on RHEL 7.5

  — Pivotal Cloud Foundry and Cloud Foundry on ESX 6.x

- In this release, the maximum supported number of PBR based external services is 200 VIPs. Scalability is expected to increase in upcoming releases.

  NOTE: With OpenShift master nodes and router nodes will be tainted by default and you might see lower scale than an upstream Kubernetes install on the same hardware.

- Cisco APIC Kubernetes and OpenShift plugins do not support the Cisco ACI Multi-Site at this time.

Caveats

This section contains lists of open and resolved caveats and known behaviors.

- Open Caveats
- Resolved Caveats

Open Caveats

This section lists the open caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Open Caveats in the 3.2(2) Release

The following are open caveats in the 3.2(2) release.

Table 4 Open Caveats in the 3.2(2) Release

Bug ID	Description	Fixed In
CSCvi09507	It is recommended that when using ESX nested Kubernetes hosts, for best performance you should provision one Kubernetes host per Kubernetes cluster on a specific ESX server.	

Resolved Caveats

This section lists the resolved caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Resolved Caveats in the 3.2(2) Release

The following table lists the resolved caveats in the 3.2(2) release.

Table 5 Resolved Caveats in the 3.2(2) Release

Bug ID	Description
CSCvj95561	The aci-container-controller of cluster X is deleting objects created by a different Kubernetes cluster in tenant common.
CSCvj90240	L3 out external network EPG lost association with contract created by plugin.
CSCvj74936	agent-ovs loses connectivity to leaf when a large SG (~250 SGRs) gets deployed.
CSCvi55118	Distributed dhcp ignores extra-dhcp-opt extensions for MTU.

Known Behaviors

This section lists caveats that describe known behaviors. Click the Bug ID to access the Bug Search Tool and see additional information about the bug.

Known Behaviors in the 3.2(2) Release

There are no known behaviors in the 3.2(2) release.

Related Documentation

The Cisco Application Policy Infrastructure Controller (APIC) documentation can be accessed from the following website:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

The documentation includes installation, upgrade, configuration, programming, and troubleshooting guides, technical references, release notes, and knowledge base (KB) articles, as well as other documentation. KB articles provide information about a specific use case or a specific topic.

By using the "Choose a topic" and "Choose a document type" fields of the Cisco APIC documentation website, you can narrow down the displayed documentation list to make it easier to find the desired document.

New Documentation

There are no new Cisco APIC product documents for this release.